Earthquake Early Warning in Eastern California

Final Technical Report - Award #: G16AC000358

8/15/2016-8/14/2017

PIs: Graham Kent and Ken Smith

gkent@unr.edu; ken@unr.edu

Nevada Seismological Laboratory, University of Nevada, Reno

MS#174, Reno NV 89511

Abstract: The Nevada Seismological Laboratory (NSL) working with partners at Caltech, UC Berkeley and the USGS focused on two areas of earthquake early warning deployment in eastern California to include: (1) installation of 5 upgraded, digital strong motion sensors (Obsidian) north and south of Lake Tahoe, and (2) understanding and reducing latencies associated with delivery of digital seismic and strong motion sensors in eastern California and western Nevada. To this end, strong motions stations (configured with L-4 seismic sensors, a marked improvement over prior L-4 analog stations) were installed at Station A (Sugar Bowl Ski Resort), Sagehen V (UC Berekely's Sagehen Field Station), Babbitt Peak, Sonora Junction and a temporary deployment near Bodie State Park. NSL also endeavored to understand latencies throughout the microwave network up and until the handoff or exchange of seismic data through either an Antelope-to-Antelope ORB2ORB process or using a SEEDLink server with a 1-second packet flush. Not surprisingly, data delivery back through the microwave network is lightning fast, on order a few 10s of milliseconds, and perhaps another half-second of in-lab processing time before the exchange to California partners. For those dataloggers with low latency delivery, such as the Obsidians that are operating as part of this project, data exiting the datalogger onboard ORB process and through to the public Internet to partners accounts for ~170 milliseconds of packet travel time, +/- a few milliseconds. Q330 datalogger deliveries were about 1.5-2.0 seconds (mostly due to Q330 1-second buffers) although a few "odd ball" dataloggers were a second or so further delayed upon that baseline (we are still investigating this issue). Tools to investigate the latencies were also produced and in the second phase of this project will be made available to earthquake early warning partners.

Report Task #1: Tahoe–Truckee Area Sensor Upgrades and Managed Wireless Microwave Communications: Efforts in Year 1 to install 5 EEW strong motion station focused in the Truckee region north of Lake Tahoe (Figure 1) and the region near Sonora Pass (Figure 2), an hour or so south of Lake Tahoe along HWY 395. Two proposed analog stations set for upgrades were abandoned to allow easier access to stations in the winter and for better ground conditions; these stations include moving analog station Independence (NN_IND) several hundred meters to Sagehen V (Fig. 3, a communication site on the UC Berkeley Sagehen research field station). Analog station Tinkers Knob (NN_TNK) with poor coupling issues and no winter access was moved to the Sugar Bowl ski resort (Station A) to improve inadequacies at Tinkers Knob and to take advantage of a significant investment of communication infrastructure as part of an extreme weather station project, now referred as Station A. In this region, the analog station at Babbitt Peak was upgraded to digital strong motion, plus an L-4 seismometer (Fig. 4, all L-4s were provided by NSL). All EEW stations installed this year include this Obsidian strong motion and L-4 seismometer pairing, except temporary site Queen Bee near Bodie, CA.

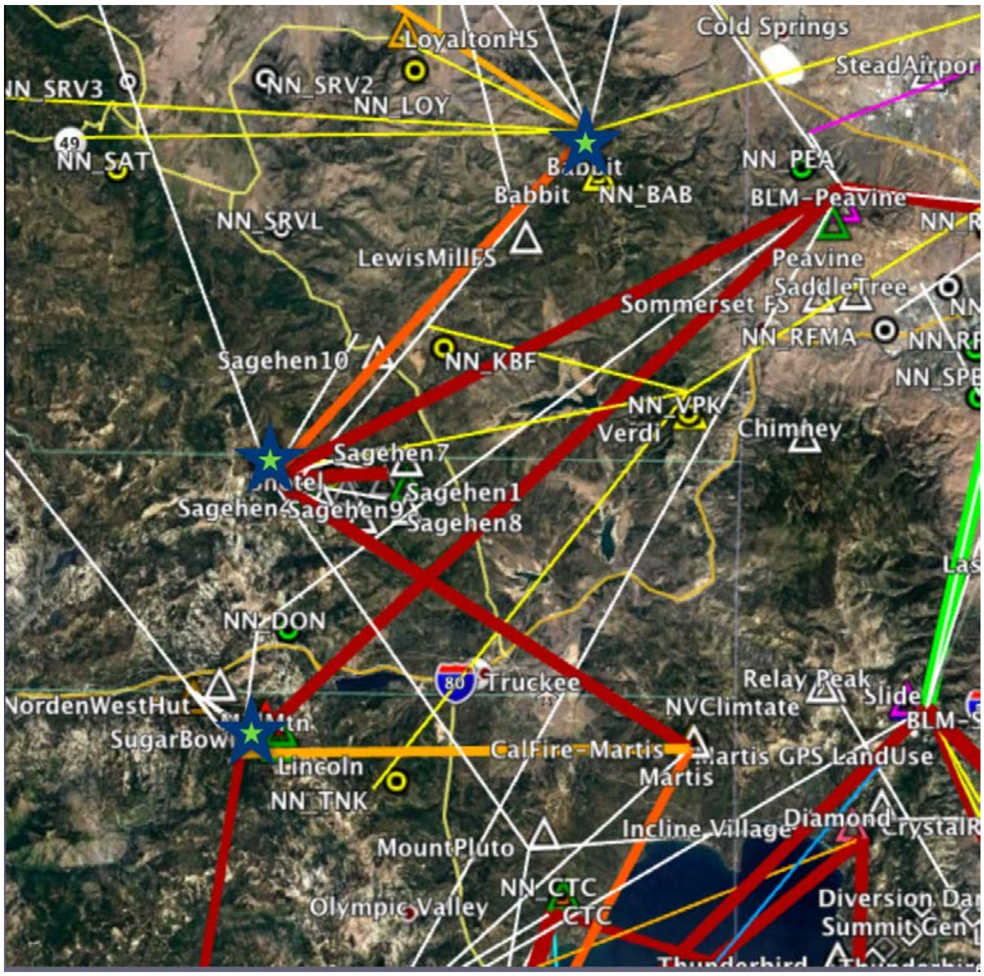

Figure 1 *Three upgraded strong motion with L-4 seismometer stations north of Lake Tahoe (see Green-Blue stars).*

South of Lake Tahoe, our Sonora Junction analog was upgraded to digital strong motion (Obsidian) with an L-4 seismometer. Unlike the first 3 EEW installs that deliver data through NSL's microwave communication system, Sonora Junction is temporarily on a cell router (with the cell tower a few hundred feet away) until our Bald Mt. communication site is upgraded to digital (funded through AlertTahoe fire program and slated for late spring 2018 install).

Figure 2 *Two upgraded strong motion with L-4 seismometer stations south of Lake Tahoe near Sonora Pass (see Green-Blue stars).*

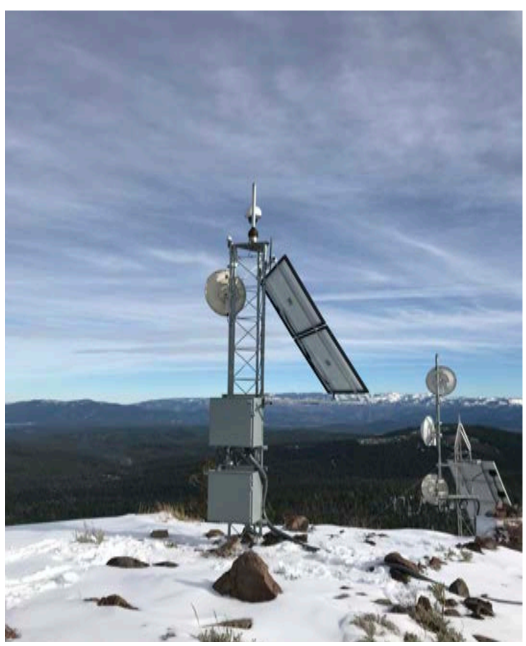

Figure 3 *Example Station at Sagehen V, a UC Berkeley Field Station north of Truckee, CA.*

Figure 4 *Example EEW vault at Station A Sugar Bowl Ski Resort, CA. Both Obsidian and L-4 seismometer are visible.*

Lastly, through working with the Tahoe National Forest and a slew of inconsistent ArcGIS land ownership maps, it was determined that our Loyalton analog site north of Truckee, CA (NN_LOY) was more likely than not on private land. All parties concerned decided it best to work with the different private landowner nearby to permit the replacement station (this problem is more common than not given upgraded GIS stations—even the Federal Government at times is unsure of ownership near boundaries). Also, NN_LOY was original sited using 1:24K maps and bruntons. As this was determined late this summer/early fall, it was decided to place our 5th Obsidian sensor and data logger at Queen Bee near Bodie, CA to take advantage of land use stemming from a fire camera to be deployed overlooking Bodie and Mono Lake/Sierra Front. Some communications infrastructure has been installed prior to camera installation. This site is not one of the 12 that UNR is presently upgrading, but remains empty, so we figured better to have the sensor (temporarily) in the field, and producing data, than on a shelf at UNR.

Task #2: Reducing and Monitoring Latency from NSL to ShakeAlert System:

Much progress has been made addressing latencies with respect to data exchange between NSL and other California ShakeAlert partners. The largest problems were rooted in SEEDLink server latencies given its current default architecture. To overcome this inadequacy, in year 1, we implemented 2 different ways to rapidly export seismic data: (1) an Antelope ORB2ORB exchange was set up between Caltech and NSL and confirmed that low latency data transfers (~ 200-350 ms) can be established amongst ShakeAlert partners over the fiber-based public Internet, and (2) re-implemented a SEEDLink server exchange with a 1 second force-flush of data buffers (in Year 2 we are working with Doug Neuhauser on sub-second SEEDlink server flush protocols). Nonetheless, there is still "jitter" within the signal that should be quantified and explained, and "can we do even better?" Although the EEW Technical Implementation Plan (OFR 2014-1097) contains no quantitative mention of network performance, at NSL we have developed preliminary tools to enable a much better understanding of internal latencies for this and other projects. A proposed objective is to more effectively work with the ShakeAlert team in incorporate contributions from NSL developments. We currently calculate the following measurements, which we believe are necessary for basic monitoring and understanding (the term "packet" refers to application-level data packets exchanged by dataloggers and our various systems):

- *Time to last packet ("downtime"):* This is a duration measurement from "now" since our datacenter last saw a packet from a given station or datalogger. This is useful for determining whether a station is operating, and is an excellent "canary" statistic for identifying problems. This metric is measured in real-time per packet.

- *Last data packet latency ("latency"):* This is a duration measurement, usually taken from the time of last data sample until the packet arrived at the data center. The specific measurement depends on the datalogger model and may differ slightly depending on acquisition specifics. This is roughly a measurement of packet travel time and processing time. This metric is measured in real-time per packet.

- *Packet latency history:* We currently run statistics on the "last packet latency" metric, using 10 second bins. We calculate low, mean, 90th percentile, and high values for each bin in real-time and store these in a time-series database for 7 days. This is extremely useful for identifying lower-level network interference, packet loss, and bandwidth issues, and is vital to visualizing the actual performance of a link/station at the application level (See Fig. 5). This is basically a mandatory prerequisite for establishing/monitoring any Quality-of-Service features at the network level, such as bandwidth reservation or traffic prioritization. For dataloggers with high transmission rates (e.g., Obsidians at 30-40pps), the latency of any one datalogger packet becomes less useful, while aggregate statistics give a much better picture of system performance.

- *Packet count ("rate"):* This is a count of packets seen per unit time. This can be usefully considered a proxy for bandwidth at a lower network layer, but can also be useful for determining the performance of a system at the application level, and becomes more important as packets become shorter in time duration and increase in number.

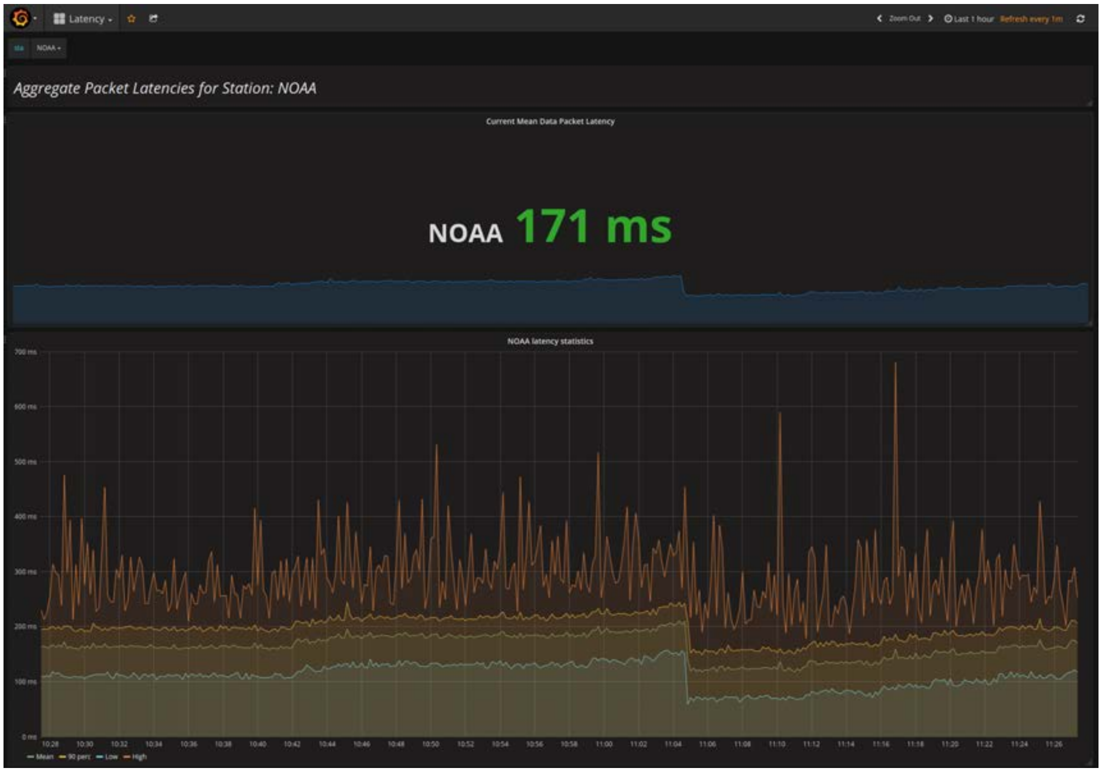

Figure 5. *Screenshot of last hour latency statistics at station NOAA (10sec bins) [low, mean, 90 percentile, high]. Large green number is the last 10s mean recorded in this time period.*

In addition to our time-series dashboards and databases, we also maintain an in-memory database for our custom "network" dashboards, so the latest metrics are always immediately available for every station (and/or channel). We expect that improved communications with ShakeAlert developers will improve all aspects of understanding latencies. NSL currently builds operational dashboard & metric systems for go/no-go active source deployments that require metric updates every 1 sec. These involve monitoring dataloggers at the channel- and configuration-settings level for changes made by field engineers (i.e., gain, sample rate changes) in real time. Our newest dashboard for earthquake station monitoring is still in development, but is attached here (see Fig. 6) for reference.

Figure 6. *Screenshot of new dashboard concept for latency monitoring, both station downtime and data packet latency. Contains link to historical statistical rollups of time series for each station. Station name is color-coded based on uptime, background is color-coded based on packet latency. For example, stations with green backgrounds are sending packets to our datacenter within 1 second.*

All of these visuals are simple wrappers around data sources. We believe that access to these data in a machine-readable format are vital to any system, especially in EEW, where a "decision module" would need to know which stations are available without human intervention. We would encourage the standardization of not only these specific metrics but more importantly, an effective, modern, secure way to share them with the downstream client (e.g., an EEW datacenter) in coordination with ShakeAlert developers. Rather than treating chunks of metric time-series data as seismic data, in the future, complete monitoring systems will most likely treat seismic data (e.g., acceleration) as another metric.

Performance

Typical packet latencies from various dataloggers: This of course depends on telemetry link quality, but these are what we normally see. The "data latency" from any given sample inside the packet will also depend on the packet length, which varies; regardless, the travel-time latencies remain fairly consistent per datalogger model:

KMI Obsidian: 100-200ms; RefTek RT130: 500ms-3s; KMI Q330S: 1-2s

Measurement Precision

With the installation of 14 Obsidians in the past year (9 Reno/Carson City ANSS SM upgrades; 5 EEW eastern California upgrades), we now measure data latencies on the order of 100 ms and lower. In order to properly measure travel times between different machines at this level of precision, NTP times are not necessarily accurate. We have preliminarily installed a Stratum 1 GPS clock-based PTP server in our data center, and are in the process of moving our monitoring and acquisition servers to this new standard. This should allow us to compare our datacenter times with the GPS times of the dataloggers at millisecond to microsecond precision. As seismic data packets become smaller and packet rates increase, there also becomes a need to scale to multiple machines within a datacenter, which will have microsecond level RPC call times. This new level of timing precision will also allow us to effectively measure latencies within and between our own datacenters, with confidence. We propose the use of funds to continue to test and implement this standard, which will be necessary for proper system design and tuning, not only for EEW but for any environmental monitoring systems in the future.